# Fast construction of the exchange operator in an atom-centered basis with concentric atomic density fitting

David S. Hollman,[1,2] Henry F. Schaefer,[1] and Edward F. Valeev[2]
[1)] *Center for Computational Quantum Chemistry, University of Georgia, 1004 Cedar St., Athens, Georgia, 30602, USA*
[2)] *Department of Chemistry, Virginia Tech, Blacksburg, Virginia 24061, USA*

A linear-scaling algorithm is presented for computing the Hartree–Fock (HF) exchange matrix using concentric atomic density fitting. The algorithm utilizes the stronger distance dependence of the three-center electron repulsion integrals along with the rapid decay of the density matrix to accelerate the construction of the exchange matrix. The new algorithm is tested with computations on systems with up to 1536 atoms and 15585 basis functions, the latter of which represents, to our knowledge, the largest quadruple-zeta HF computation ever performed. Our method handles screening of high angular momentum contributions in a particularly efficient manner, allowing the use of larger basis sets for large molecules without a prohibitive increase in cost.

arXiv:1410.4882v1 [physics.chem-ph] 17 Oct 2014

## I. INTRODUCTION

Computation of the so-called exchange operator is often the most expensive step in electronic structure models applicable to large systems, such as the hybrid[1] Density Functional Theory (DFT) and Hartree-Fock (HF),[2] the latter of which serves as the starting point for electron correlation treatment with modern reduced-scaling many-body methods.[3] The exchange matrix, $\mathbf{K}$, is given by

$$K_{\mu\nu} = \sum_{\lambda\sigma} D^{\lambda\sigma} \left(\mu\lambda|\nu\sigma\right), \tag{1}$$

where $\mu$, $\nu$, $\lambda$, $\sigma$, ... denote basis functions, $\mathbf{D}$ is the density matrix, and

$$(\mu\lambda|\nu\sigma) = \int \chi_\mu(\vec{r}_1)\chi_\lambda(\vec{r}_1)\frac{1}{|\vec{r}_1 - \vec{r}_2|}\chi_\nu(\vec{r}_2)\chi_\sigma(\vec{r}_2)d\tau_1 d\tau_2 \tag{2}$$

are the electron repulsion integrals (ERIs) in the Mulliken "bra-ket" notation. Here we consider atom-centered basis functions $\chi(\vec{r})$, typically represented by (contracted) Gaussian-type orbitals. Although $(\mu\lambda|\nu\sigma)$ decays with the bra-ket distance $R$ only as $R^{-1}$ (i.e., slowly), the decay of the density matrix $D^{\lambda\sigma}$ in a finite system (and, generally, any system with a nonzero band gap) is exponential with the distance between $\lambda$ and $\sigma$.[4] Therefore, as first noted by Almlöf,[5] the number of nontrivial elements of the exchange matrix $\mathbf{K}$ grows linearly with system size—that is, $\mathbf{K}$ has $O(N)$ significant elements ($N \propto$ system size). A number of linear-scaling methods for the computation of $\mathbf{K}$ have been investigated over the years, with pioneering work done in chemistry in the mid-1990s by Schwegler, Challacombe, and Head-Gordon[6,7] and Burant, Scuseria, and Frisch.[8] Several related algorithms have been developed, e.g. ONX by Challacombe et al and `LinK` by Ochsenfeld, White, and Head-Gordon[9].

The aforementioned $O(N)$ algorithms for exchange take advantage of the element-wise sparsity of the ERI and density matrix tensors in Eq. (1). Some effort has gone into utilization of rank sparsity of the ERI tensor, as revealed by the multipole expansion of the Coulomb operator,[10] by the pseudospectral and related factorizations of ERI,[11–15] or, as done here, by the density fitting approximation of ERIs.[16] The density fitting approximation[17–26] (DF, also called resolution of the identity), expresses the two-center products in the bra and the ket of the ERI tensor (often called "densities," not to be confused with the density matrix $\mathbf{D}$) as linear combinations of one-center functions from an auxiliary basis:

$$|\mu\nu) = \sum_X C^X_{\mu\nu} |X), \tag{3}$$

where $X$, $Y$, ... are indices in an auxiliary basis set. To obtain the coefficients $C^X_{\mu\nu}$, one must left-project with the auxiliary basis to obtain a system of linear equations:

$$(Y|\mu\nu) = \sum_X C^X_{\mu\nu} (Y|X). \tag{4}$$

The solution of these equations in their full form requires $O(N^3)$ effort and is thus untenable for large molecular systems. Moreover, substitution of the form in Eq. (3) into the exchange matrix expression in Eq. (1) does not reduce the formal scaling of the exchange matrix build. Consequently, significant work has been done on methods of reducing the scaling of Eq. (3), principally by localizing the fitting basis.[16,19,27–34] Significantly less research has been done on the use of these localized decompositions for the exchange matrix build, Eq. (1), particularly in the context of the linear scaling exchange work of the late 1990s. Indeed, only Sodt and Head-Gordon[16] have investigated a linear scaling exchange construction with local DF. Also, Neese and coworkers[12] presented a linear scaling algorithm for HF exchange in the context of their "chain-of-spheres" exchange (COSX) method, which borrows elements from both DF and the pseudo-spectral method of Friesner, et al.[11] Herein, we present a linear scaling HF exchange algorithm and implementation using the Concentric Atomic Density Fitting (CADF) ap-

proach, based on a concept that has been around for some time[19,30] and was recently revived by Merlot, et al.[34] and the present authors.[33]

In Section II, we briefly review the theory behind CADF and develop a diagrammatic scheme for further discussion of linear scaling methods. Section III presents a progression of algorithms that lead to the final CADF-LinK method, and Section IV discusses implementation details. Section V gives the results of some benchmark computations, and Section VI contains some concluding remarks.

## II. THEORETICAL BACKGROUND

### A. Concentric Atomic Density Fitting

The idea behind what we call CADF has been around since the early days of DF:[19] expand a product density $|\mu\nu)$ using only auxiliary functions centered on the same atom as the composite functions $\mu$ and $\nu$. The fitting equations can thus be written

$$\left(Y_{(ab)}\middle|\mu_a\nu_b\right) = \sum_{X\in(ab)} C^{X_{(ab)}}_{\mu_a\nu_b} \left(X_{(ab)}\middle|Y_{(ab)}\right), \qquad (5)$$

where $a$, $b$, ... are indices of atom centers, $\mu_a$ indicates that the function $\mu$ is centered on atom $a$, and $Y_{(ab)}$ indicates that the function $Y$ is centered on either atom $a$ or atom $b$. In this context, the use of the simple expansion of the integral tensor

$$\begin{aligned} g_{\mu\nu,\lambda\sigma} &\equiv (\mu_a\nu_b|\lambda_c\sigma_d) \\ &\approx \sum_{X\in(ab)} \sum_{Y\in(cd)} C^{X_{(ab)}}_{\mu_a\nu_b} \left(X_{(ab)}\middle|Y_{(cd)}\right) C^{Y_{(cd)}}_{\lambda_c\sigma_d} \end{aligned} \qquad (6)$$

yields errors in $\mathbf{g}$ that are too large to be useful for HF theory computations. However, in 2000 Dunlap[35] noted that this expansion is linear in the density error (unless the full, non-local Coulomb metric is used), and a more robust expansion correct to second order in the density error can be constructed. In the context of CADF, this robust expansion can be written

$$\begin{aligned} (\mu_a\nu_b|\lambda_c\sigma_d) \approx & \sum_{X\in(ab)} C^{X_{(ab)}}_{\mu_a\nu_b} \left(X_{(ab)}\middle|\lambda_c\sigma_d\right) \\ & + \sum_{Y\in(cd)} \left(\mu_a\nu_b\middle|Y_{(cd)}\right) C^{Y_{(cd)}}_{\lambda_c\sigma_d} \\ & - \sum_{X\in(ab)} \sum_{Y\in(cd)} C^{X_{(ab)}}_{\mu_a\nu_b} \left(X_{(ab)}\middle|Y_{(cd)}\right) C^{Y_{(cd)}}_{\lambda_c\sigma_d} \end{aligned} \qquad (7)$$

The use of this robust expansion yields reasonable errors for HF theory—on the order of 2-5 times the errors arising from conventional density fitting, in most cases.[33,34] However, Eq. (7) breaks the positive semidefiniteness of the $\mathbf{g}$ tensor,[34] which can cause convergence issues in rare cases. Merlot, et al. and our group both suggested different approaches[33,34] to correcting this issue. While we consider the convergence issues with CADF to be eminently solvable, to throughly test the CADF approximation we need to be able to apply CADF-based Hartree-Fock and Density Functional Theory methods to large molecules. Thus, in this paper the focus is on developing a practical $O(N)$ CADF-based exchange matrix algorithm, and hence the issues of positivity are set aside herein.

### B. Screening of Insignificant Contributions

Due to the rapid decay of the overlap between the Gaussian basis function functions $\mu$ and $\nu$, only $O(N)$ products $|\mu\nu)$ result in significant ERIs. Thus, the four-center ERI tensor $\mathbf{g}$ contains $O(N^2)$ significant entries. Häser and Ahlrichs[36] noted that the Cauchy–Schwarz inequality holds for the ERI tensor, and thus this exponential decay could be exploited for prescreening using the inequality

$$|\,(\mu\nu|\lambda\sigma)\,| \le |\,(\mu\nu|\mu\nu)\,|^{1/2}\,|\,(\lambda\sigma|\lambda\sigma)\,|^{1/2}. \qquad (8)$$

Using the nomenclature of Neese and coworkers,[12] we say that $\mu$ forms an "S-junction" with $\nu$ if and only if $|\,(\mu\nu|\mu\nu)\,|^{1/2}$ is greater than some pairing threshold $\epsilon_S$.

#### 1. Diagrammatic Notation

To illustrate the screening of significant contributions to the exchange matrix, we introduce a simple graphical method for showing contributing factors to the sparsity of a given tensor or tensor contraction expression. Tensor indices are shown as vertices of a diagram. Each of these indices has a range of $O(N)$, hence a diagram with $k$ vertices denotes to $O(N^k)$ total contributions to a tensor or a tensor contraction. An edge connecting two vertices denotes that only $O(N)$ *pairs* of these indices are significant for any finite precision in the limit of infinite system size. In other words, for each first index value there is a $O(1)$ list of significant second index values. For example, the ERI tensor screened using Eq. (8) can be represented as

$$(\mu\lambda|\nu\sigma) \;:\quad \mu \overset{\text{S}}{\text{———}} \lambda \qquad\qquad \nu \overset{\text{S}}{\text{———}} \sigma \qquad (9)$$

From this diagram, it is immediately apparent that the screening in Eq. (8) requires $O(N^2)$ integrals to be computed for $\mathbf{g}$, since there are a constant number of $\lambda$ for each $\mu$ and a constant number of $\sigma$ for each $\nu$, but no connection between $\mu$ and $\nu$. Therefore, the diagram must be fully connected for a tensor or contraction to have $O(N)$ significant contributions.

### 2. Linear Scaling Exchange

$O(N)$ exchange construction algorithms take advantage of the aforementioned rapid decay of the density matrix.[6–9] Diagrammatically, the conventional exchange matrix expression

$$K_{\mu\nu} = \sum_{\lambda\sigma} D^{\lambda\sigma} \, (\mu\lambda|\nu\sigma) \tag{10}$$

can be represented as

$$K_{\mu\nu} \;:\quad \mu \overset{\text{S}}{\text{———}} \lambda \overset{\text{P}}{\text{———}} \sigma \overset{\text{S}}{\text{———}} \nu \tag{11}$$

where the connection between $\lambda$ and $\sigma$ is called a "P-junction" in the nomenclature of Neese, et al.[12] Diagram (11) is fully connected, which suggests the existence of a linear scaling algorithm. Note also that the path length between two indices in the result suggests (in a qualitative sense) some aspects of the performance of the associated linear scaling algorithm. Since there are a bounded number of $\lambda$ for a given $\mu$—we will use the notation $C_S(\mu)$—as well as a bounded number of $\sigma$ given $\lambda$ and a bounded number of $\nu$ given $\sigma$, the prefactor of the linear scaling algorithm will be bounded from above by the product $C_S(\mu)C_P(\lambda)C_S(\sigma)$, because one could, at worst, use the direct product of these index sets to construct the tensor. In practice, the scaling will be much better than this (since, e.g., the product of medium-sized contributions to each connection may still be quite small), though there may not necessarily be a reasonable algorithm to access this scaling. Nevertheless, the diagrammatic approach presents a concise picture of the nature of the factors that necessarily contribute to the performance of any linear scaling algorithm for the construction of **K**. It is clear from the form of the Gaussian product rule (GPR) that $C_S(\mu)$ and $C_S(\sigma)$ will depend on the number of basis functions per atom and the diffuseness of those functions. From physical reasoning, we note that $C_P(\lambda)$ will further depend on the band gap of the system and the degree of delocalization of electrons in the system. Given these properties of S and P junctions, we conclude that a linear scaling algorithm associated with this diagram would probably perform poorly for large basis sets or systems with small band gaps. While this is not a particularly profound insight (both of these are well-known issues with virtually all linear scaling methods), it suggests by contrast what a method that improves on these limitations might look like: such a method would have to provide an alternative path between $\mu$ and $\nu$ that does not rely on the square of the S junction constant (for large basis sets) or the P-junction constant (for smaller band-gap systems). Our task herein is to identify useful additional pathways between $\mu$ and $\nu$.

### 3. CADF with Linear Scaling Exchange

CADF introduces another kind of junction between indices, that of concentricity. Using the diagrammatic approach, we can write the coefficient tensor as a union of two diagrams:

$$C^{X_{(ab)}}_{\mu_a\lambda_b} \;:\quad \begin{array}{l} \mu_a \overset{\text{S}}{\text{———}} \lambda_b \\ \text{CA}\big| \\ X_a \end{array} \quad \cup \quad \begin{array}{l} \mu_a \overset{\text{S}}{\text{———}} \lambda_b \\ \text{CA}\big| \\ X_b \end{array} \quad = \quad \begin{array}{l} \mu_a \overset{\text{S}}{\text{———}} \lambda_b \\ \text{CA}\,\sim\!\sim\!\sim \\ X_{(ab)} \end{array} \tag{12}$$

where "CA" is short for "Concentric Atomic," indicating that the connected indices are centered on the same atom, and the squiggled lines are used to introduce a shorthand for the left-hand side of the equals sign. By plugging the robust CADF approximation to ERI, Eq. (7), into Eq. (10) and relabeling some indices, we can write **K** as

$$K_{\mu_a\nu_c} = \tilde{K}_{\mu_a\nu_c} + \tilde{K}_{\nu_c\mu_a}, \tag{13}$$

$$\tilde{K}_{\mu_a\nu_c} = \sum_{\sigma_d X_{(cd)}} \sum_{\lambda_b} \left[ \left(\mu_a\lambda_b \middle| X_{(cd)}\right) - \tfrac{1}{2} \sum_{Y_{(ab)}} C^{Y_{(ab)}}_{\mu_a\lambda_b} \left(Y_{(ab)} \middle| X_{(cd)}\right) \right] D^{\lambda_b\sigma_d} C^{X_{(cd)}}_{\nu_c\sigma_d}. \tag{14}$$

Diagrammatically, this can be written

$$\tilde{K}_{\mu_a\nu_c} \;:$$ 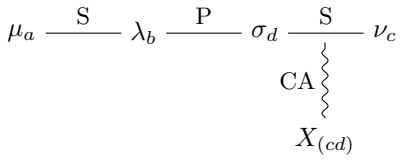

$$-\;\frac{1}{2}$$ 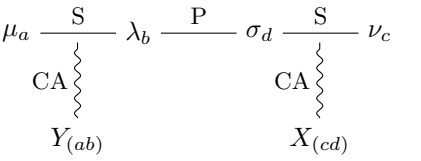

(15)

Both terms in the diagram are connected, hence a linear-scaling algorithm can be designed. However, it is possible to significantly reduce the computational cost of the first term by screening the three-center two-electron integrals.

#### 4. Three-Center Integral Screening with SQV$\ell$

Three-center ERIs can be screened more efficiently than the standard four-center ERIs because the potential created by a solid-harmonic Gaussian $|X)$ of angular momentum $l_X$ decays as $R^{-(l_X+1)}$, rather than the $R^{-1}$ decay of a general product $|\nu\sigma)$. The decay of the Coulomb operator does not matter as much for the four-center ERI-based exchange construction because the density decays so rapidly. In the CADF-based approach, efficient screening of the three-center ERIs is important and cannot be achieved using only the Schwartz bound. Recently we developed a tight estimator, called SQV$\ell$, for three-center Coulomb integrals that takes into account the correct asymptotic decay of the potential:[37]

$$(\kappa\eta|\theta) \approx \begin{cases} (2\pi)^{3/4}\,\beta_{\ell_\theta}(\zeta_\theta)\dfrac{|S_{\kappa\eta}|}{R^{\ell_\theta+1}} & R > \mathrm{ext}_{\kappa\eta}+\mathrm{ext}_\theta \\ & \text{and } S_{\kappa\eta}/Q_{\kappa\eta} > \vartheta_{\mathrm{SQ}} \\ \dfrac{\pi\sqrt{2}\,\beta_{\ell_\theta}(\zeta_\theta)\,Q_{\kappa\eta}}{(\zeta_\kappa+\zeta_\eta)^{\frac{1}{4}}R^{\ell_\theta+1}} & R > \mathrm{ext}_{\kappa\eta}+\mathrm{ext}_\theta \\ & \text{and } S_{\kappa\eta}/Q_{\kappa\eta} \le \vartheta_{\mathrm{SQ}} \\ Q_{\kappa\eta}Q_\theta & R \le \mathrm{ext}_{\kappa\eta}+\mathrm{ext}_\theta, \end{cases} \tag{16}$$

where

$$\beta_\ell(\zeta) \equiv \zeta^{-\frac{2\ell+3}{4}}\sqrt{(2\ell-1)!!}, \tag{17}$$

$$Q_{\mu\nu} \equiv (\mu\nu|\mu\nu)^{1/2}, \tag{18}$$

$\zeta_\kappa$, $\zeta_\eta$, and $\zeta_\theta$ are the exponents of the various primitives, $S_{\kappa\eta}$ is the overlap, $\mathrm{ext}_{\kappa\eta}$ are CFMM extents, and $\vartheta_{\mathrm{ws}}$ and $\vartheta_{\mathrm{SQ}}$ are user-defined thresholds (see Ref. 37 for details, including generalization to contracted basis functions). Practically speaking, the ratio of prefactors from the first two cases is folded into the $S_{\kappa\eta}/Q_{\kappa\eta}$ ratio, and the second case is used if this ratio is less than one. Effectively, the far field estimator uses the minimum of the first two cases when the $S_{\kappa\eta}/Q_{\kappa\eta}$ ratio is greater than $\vartheta_{\mathrm{SQ}}$, and only the second case when the ratio is less than $\vartheta_{\mathrm{SQ}}$. This latter detail will be important for the use of this estimator in Section III.

The SQV$\ell$ estimator allows us to add a crucial edge to the CADF exchange screening diagram:

$$\tilde{K}_{\mu_a\nu_c} \;\leftarrow:$$ 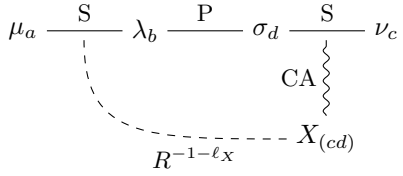

(19)

Note that the efficiency of the SQV$\ell$ screening will increase for kets with higher angular momenta; hence, the more expensive the three-center ERI, the more likely it will be screened out.

## III. $O(N)$ CADF EXCHANGE ALGORITHM

Like four-center ERI-based linear scaling exchange algorithms,[7,9] our linear scaling exchange algorithm based on concentric atomic density fitting (which we will call CADF-LinK) relies on the construction of prescreening lists, $L^{(3)}_{\mu X}$ and $L^{(B)}_{\mu X}$, to drive the exchange matrix build. Both sets of lists are computed with linear effort.

### A. The $L^{(3)}_{\mu X}$ Lists

$L^{(3)}_{\mu X}$ specifies the list of $\lambda$ for a given $(\mu X)$ pair for which $(\mu\lambda|X)$ must be computed. It is defined as follows:

$$\lambda \in L^{(3)}_{\mu X} \;\Leftrightarrow\; \bar{d}^X_\lambda \tilde{R}^X_{\mu\lambda} > \epsilon, \tag{20}$$

where

$$\bar{d}^X_\lambda \equiv \sum_\sigma^{OBS} |D^{\lambda\sigma}|\bar{C}^X_\sigma, \tag{21}$$

$$\bar{C}^{X_c}_{\sigma_d} \equiv (X|X)^{1/2} \underset{\nu_e\in OBS}{\mathrm{Frob}} \left\{C^{X_c}_{\nu_e\sigma_d}\delta_{c\in(de)}\right\} \tag{22}$$

$$= (X|X)^{1/2}\left[\underset{\nu_c\in(c)_{OBS}}{\mathrm{Frob}}\left\{C^{X_c}_{\nu_c\sigma_d}\right\} + \underset{\nu_e\in OBS}{\mathrm{Frob}}\left\{C^{X_c}_{\nu_e\sigma_d}\right\}\delta_{cd} - \underset{\nu_c\in(c)_{OBS}}{\mathrm{Frob}}\left\{C^{X_c}_{\nu_c\sigma_d}\right\}\delta_{cd}\right], \tag{23}$$

$$\tilde{R}^X_{\mu\lambda} \equiv (X|X)^{-1/2}\, I_{\mathrm{SQV}\ell}(\mu,\lambda,X) \tag{24}$$

and $I_{\mathrm{SQV}\ell}(\mu,\lambda,X)$ is the SQV$\ell$ screening factor from Eq. (16). (Keep in mind that though the tensor-like notation $\tilde{R}^X_{\mu\lambda}$ is used for compactness, this quantity is intended to be computed on the fly, so a notation like $\tilde{R}(\mu,\lambda,X)$ would be conceptually more appropriate.) We have loosely used the notation

$$\underset{x\in\mathbb{S}}{\mathrm{Frob}}\,\{f_x\} = \sqrt{\sum_{x\in\mathbb{S}} f_x^2} \tag{25}$$

to represent the block Frobenius norm of a matrix or tensor expression. In these expressions and the following, the orbital basis set (OBS) and density fitting basis set (DFBS) are distinguished by their respective abbreviations, and the notation $(c)_{OBS}$, for instance, represents the set of all basis functions from the OBS centered on atom $c$. The density fitting coefficients $C^{X_c}_{\nu_e\sigma_d}$ can be obtained in linear effort after a Schwarz-based prescreening of OBS pairs, as discussed in the original CADF paper.[33]

We immediately note that (theoretically, at least) the $\mathbf{đ}$ ("d-bar") intermediate can be constructed in linear scaling time, since

$$đ^X_\lambda \;:\; \lambda_b \overset{\mathrm{P}}{\text{———}} \sigma_d \overset{\mathrm{S}}{\text{———}} \nu_c \;\; \overset{\mathrm{CA}}{\sim} \; X_{(cd)} \tag{26}$$

is a fully connected diagram. Indeed, the construction of $\mathbf{đ}$ is quite straightforward if the density matrix $\mathbf{D}$ is stored in a sparse data structure. The construction of $\bar{\mathbf{C}}$ can also be performed in linear scaling effort, since the norm in the second term of Eq. (23) need only loop over the $\nu_e$ that have significant overlap with $\sigma_d$. With this in mind, a procedure for the linear scaling construction of $\mathbb{L}^{(3)}$ is given in Algorithm 1. Similar to the original `LinK` algorithm,[9] the key to linear scaling in Algorithm 1 is the list ordering that allows the program to exit the loops early in lines 12 and 16. The auxiliary list set $L^{(DC)}_\lambda$ is composed of $X$ for which $đ^X_\lambda$ is significant; each list is then sorted by decreasing $đ^X_\lambda$ value. The other helper list set, $L^{(S)}_\lambda$, is simply the $\mu$ that form Schwarz pairs with $\lambda$, sorted by decreasing value of $|\,(\mu\lambda|\mu\lambda)\,|^{1/2}$.

Algorithm 1: Build $\mathbb{L}^{(3)}$

1: **foreach** $\lambda$ **in** $OBS$ **do**
2: $\quad L^{(DC)}_\lambda \leftarrow X \in DFBS$ with $đ^X_\lambda > \epsilon_d$
3: $\quad$ Sort $L^{(DC)}_\lambda$ by decreasing $|đ^X_\lambda|$
4: $\quad$ **foreach** $X$ **in** $L^{(DC)}_\lambda$ **do**
5: $\quad\quad$ **foreach** $\mu$ **in** $L^{(S)}_\lambda$ **do**
6: $\quad\quad\quad$ **if** $đ^X_\lambda\,(\mu\lambda|\mu\lambda)^{1/2} > \epsilon_K$ **then**
7: $\quad\quad\quad\quad$ Compute distance factor $\tilde{R}^X_{\mu\lambda}$
8: $\quad\quad\quad\quad$ **if** $đ^X_\lambda \tilde{R}^X_{\mu\lambda} > \epsilon_K$ **then**
9: $\quad\quad\quad\quad\quad$ Add $\lambda$ to $L^{(3)}_{\mu X}$
10: $\quad\quad\quad\quad$ **end if**
11: $\quad\quad\quad$ **else**
12: $\quad\quad\quad\quad$ **break** out of $\mu$ loop
13: $\quad\quad\quad$ **end if**
14: $\quad\quad$ **end** loop over $\mu$
15: $\quad\quad$ **if** $\mu$ loop exited on first iteration **then**
16: $\quad\quad\quad$ **break** out of $X$ loop
17: $\quad\quad$ **end if**
18: $\quad$ **end** loop over $X$
19: **end** loop over $\lambda$

### 1. Quadratic Exchange Matrix Build

Using only the $\mathbb{L}^{(3)}$ lists, a quadratic scaling $\mathbf{K}$ build algorithm can be devised. Such a procedure is outlined in Algorithm 2. Note that the subscripted indices (e.g., $\mu_a$) are descriptive rather than prescriptive: the subscript $a$ merely indicates that the center on which the function $\mu$ is located has the index $a$; it does not indicate anything specific about the iteration range of $\mu_a$. The algorithm loops over the significant $(\mu\lambda|X)$ integrals and immediately contracts these contributions into the $\mathbf{B}$ intermediate. The quadratic scaling of this algorithm arises because the loop in line 10 is not restricted; the restriction of this loop requires $\mathbb{L}^{(B)}$, discussed in Section III B.

Several important aspects of this quadratic scaling algorithm merit further discussion before the introduction of the linear scaling version. First, note that the summation of $\mathbf{B}$ into $\tilde{\mathbf{K}}$ (lines 13–22 of Algorithm 2) is split into two loops with one index restricted in each loop. This restriction is made possible by the CADF coefficient definition, which precludes coefficients with DFBS functions that do not share a center with the OBS pair being fit. This loop structure also makes it possible to store the coefficients in an ordering more efficient for vectorization: $C^{X_c}_{\nu_c\sigma_d}$ is stored as $X_c \rightarrow [\nu_c \text{ offset in } (c)_{OBS}] \rightarrow \sigma_d$ (from slowest- to fastest-running index), which is still a $O(N^2)$ data structure. Secondly, note that the algorithm is integral direct, and that the algorithm's storage requirements never exceed $O(N^2)$. The intermediate $\mathbf{B}$ requires linear storage in this algorithm, and the integrals $(\mu\lambda|X)$ need not be stored outside of the inner loop in which they are computed.

## B. The $L^{(B)}_{\mu X}$ Lists

In order to further improve the scaling of Algorithm 2 to linear, another set of prescreening lists, $\mathbb{L}^{(B)}$, must be formed. For a given $\mu$ and $X$, the list $L^{(B)}_{\mu X}$ from the set $\mathbb{L}^{(B)}$ is defined as

$$\sigma \in L^{(B)}_{\mu X} \Leftrightarrow \bar{C}^X_\sigma \bar{b}^\sigma_{(\mu X)} > \epsilon \tag{27}$$

Algorithm 2: Quadratic $\mathbf{K}$ build using $\mathbb{L}^{(3)}$

1: Form $L^{(3)}_{\mu X}$ for significant $(\mu_a X_c)$ ▷ See Algorithm 1
2: **foreach** $(\mu_a X_c)$ **with** $L^{(3)}_{\mu X} \neq \varnothing$ **do**
3: **foreach** $\lambda_b$ **in** $L^{(3)}_{\mu X}$ **do**
4: Compute $(\mu_a \lambda_b | X_c)$
5: $\bar{g}^{X_c}_{\mu_a \lambda_b} = (\mu_a \lambda_b | X_c)$
6: **foreach** $Y_{(ab)}$ **in** $(a)_{DFBS} \cup (b)_{DFBS}$ **do**
7: $\bar{g}^{X_c}_{\mu_a \lambda_b} -= \frac{1}{2} C^{Y_{(ab)}}_{\mu_a \lambda_b} \left(Y_{(ab)} \middle| X_c\right)$
8: **end** loop over $Y$
9: **foreach** $\sigma_d$ **in** OBS **do**
10: $B^{X_c}_{\mu_a \sigma_d} += \bar{g}^{X_c}_{\mu_a \lambda_b} D^{\lambda_b \sigma_d}$
11: **end** loop over $\sigma_d$
12: **end** loop over $\lambda_b$
13: **foreach** $\nu_c$ **in** $(c)_{OBS}$ **do**
14: **foreach** $\sigma_d$ **in** $L^{(S)}_{\nu}$ **with** $\sigma_d \notin (c)_{OBS}$ **do**
15: $\tilde{K}_{\mu_a \nu_c} += C^{X_c}_{\nu_c \sigma_d} B^{X_c}_{\mu_a \sigma_d}$
16: **end** loop over $\sigma_d$
17: **end** loop over $\nu_c$
18: **foreach** $\sigma_c$ **in** $(c)_{OBS}$ **do**
19: **foreach** $\nu_d$ **in** $L^{(S)}_{\sigma}$ **do**
20: $\tilde{K}_{\mu_a \nu_d} += C^{X_c}_{\sigma_c \nu_d} B^{X_c}_{\mu_a \sigma_c}$
21: **end** loop over $\nu_d$
22: **end** loop over $\sigma_c$
23: **end** loop over $(\mu_a X_c)$
24: $\mathbf{K} = \tilde{\mathbf{K}} + \tilde{\mathbf{K}}^\top$

where

$$\mathfrak{b}^{\sigma}_{(\mu X)} \equiv \sum_{\lambda \in L^{(3)}_{\mu X}} |D^{\lambda\sigma}| \tilde{R}^{X}_{\mu\lambda}. \tag{28}$$

As with $\mathfrak{d}$, the formation of $\mathfrak{b}$ in linear effort is relatively trivial if $\mathbf{D}$ is stored in a sparse data structure, and particularly if $L^{(3)}_{\mu X}$ is sorted by decreasing $\tilde{R}^{X}_{\mu\lambda}$. An algorithm for the linear scaling formation of $\mathbb{L}^{(B)}$ is outlined in Algorithm 3. The procedure is greatly simplified by the fact that a list of significant $(\mu X)$ pairs is already known from the construction of $\mathbb{L}^{(3)}$. The algorithm is little more than a sparse-sparse matrix multiply. If $\bar{\mathbf{C}}$ is stored in a sparse data structure, line 2 is already part of the manifestation of $\bar{\mathbf{C}}$ in memory. Similarly, the conditional in line 6 is essentially the procedure that a sparse-sparse matrix multiply undergoes to determine the nontrivial entries in the product matrix. Nonetheless, we have included the $\mathbb{L}^{(B)}$ build algorithm here for completeness.

### 1. Linear Scaling Exchange Matrix Build

The use of $\mathbb{L}^{(B)}$ to create a $\mathbf{K}$ build procedure that scales linearly requires only a few modifications to Algorithm 2. The revised procedure is given in Algorithm 4. The main changes involve restricting the loops in lines 10 and 15. The restriction in line 10 amounts to excluding from $\mathbf{B}$ any $\sigma_d$ where the contraction over the full set of significant $\lambda_b$ for a given $(\mu_a X_c)$ is negligible. These indices can then also be excluded from the summation in line 15, since the $B^{X_c}_{\mu_a \sigma_d}$ are trivial for these $\sigma_d$. Further optimization could be made by restricting the loop in line 10 to include only $\sigma_d$ for a given $\lambda_b$ where $D^{\lambda_b \sigma_d}$ is significant. This optimization would arise naturally if a sparse data structure were used for $\mathbf{D}$, but in the present implementation dense $O(N^2)$ data structures were used (see Section IV).

Algorithm 3: Build $\mathbb{L}^{(B)}$

1: **foreach** $X$ **in** $OBS$ **do**
2: $L^{(\bar{C})}_{X} \leftarrow \sigma \in OBS$ with $\bar{C}^{X}_{\sigma} > \epsilon_{\bar{C}}$
3: **end** loop over $\sigma$
4: **foreach** $(\mu X)$ **with** $L^{(3)}_{\mu X} \neq \varnothing$ **do**
5: **foreach** $\sigma$ **in** $L^{(\bar{C})}_{X}$ **do**
6: **if** $\bar{C}^{X}_{\sigma} \mathfrak{b}^{\sigma}_{(\mu X)} > \epsilon_K$ **then**
7: Add $\sigma$ to $L^{(B)}_{\mu X}$
8: **end if**
9: **end** loop over $\sigma$
10: **end** loop over $(\mu X)$

## IV. COMPUTATIONAL DETAILS

The algorithms detailed in Section III were implemented in a development version (commit tag `3.0.0-cadflink` of the `localdf` branch) of the Massively Parallel Quantum Chemistry (`MPQC`)[38] quantum chemistry package. While the algorithms were implemented predominantly as written here, several small details differ from the ideal linear scaling implementation. Most prominently, since integrals are computed more efficiently in shell blocks, all of the indices in the algorithms actually represent shell blocks rather than individual basis functions. In other words, for instance,

$$|\,(X|X)\,|^{1/2} = \underset{X' \in X}{\text{Frob}} \left\{ |\,(X'|X')\,|^{1/2} \right\} \tag{29}$$

where $X'$ is a function index and $X$ is a shell index. The upshot of this is that the prescreening algorithms and list formations are significantly less expensive than the main computation in practice, even though in principle the onset of linear scaling is later for these portions of the computation. (The prescreening is more expensive in principle because the exit conditions of the loops in Algorithm 1 use the Schwarz estimate of the integrals rather than the distance-including estimate, since the latter cannot be easily ordered). The quantities used in the screening process are replaced by their shell block Frobe-

Algorithm 4: Linear Scaling $\mathbf{K}$ build using $\mathbb{L}^{(3)}$ and $\mathbb{L}^{(B)}$

1: Form $L^{(3)}_{\mu X}$ for significant $(\mu_a X_c)$ ▷ See Algorithm 1
2: Form $L^{(B)}_{\mu X}$ for significant $(\mu_a X_c)$ ▷ See Algorithm 3
3: **foreach** $(\mu_a X_c)$ **with** $L^{(3)}_{\mu X} \neq \varnothing$ and $L^{(B)}_{\mu X} \neq \varnothing$ **do**
4: **foreach** $\lambda_b$ **in** $L^{(3)}_{\mu X}$ **do**
5: Compute $(\mu_a \lambda_b | X_c)$
6: $\bar{g}^{X_c}_{\mu_a \lambda_b} = (\mu_a \lambda_b | X_c)$
7: **foreach** $Y_{(ab)}$ **in** $(a)_{DFBS} \cup (b)_{DFBS}$ **do**
8: $\bar{g}^{X_c}_{\mu_a \lambda_b} -= \frac{1}{2} C^{Y_{(ab)}}_{\mu_a \lambda_b} \left( Y_{(ab)} \middle| X_c \right)$
9: **end** loop over $Y$
10: **foreach** $\sigma_d$ **in** $L^{(B)}_{\mu X}$ **do**
11: $B^{X_c}_{\mu_a \sigma_d} += \bar{g}^{X_c}_{\mu_a \lambda_b} D^{\lambda_b \sigma_d}$
12: **end** loop over $\sigma_d$
13: **end** loop over $\lambda_b$
14: **foreach** $\nu_c$ **in** $(c)_{OBS}$ **do**
15: **foreach** $\sigma_d$ **in** $L^{(B)}_{\mu X}$ **with** $\sigma_d \notin (c)_{OBS}$ **do**
16: $\tilde{K}_{\mu_a \nu_c} += C^{X_c}_{\nu_c \sigma_d} B^{X_c}_{\mu_a \sigma_d}$
17: **end** loop over $\sigma_d$
18: **end** loop over $\nu_c$
19: **foreach** $\sigma_c$ **in** $(c)_{OBS} \cap L^{(B)}_{\mu X}$ **do**
20: **foreach** $\nu_d$ **in** $L^{(S)}_{\sigma}$ **do**
21: $\tilde{K}_{\mu_a \nu_d} += C^{X_c}_{\sigma_c \nu_d} B^{X_c}_{\mu_a \sigma_c}$
22: **end** loop over $\nu_d$
23: **end** loop over $\sigma_c$
24: **end** loop over $(\mu_a X_c)$
25: $\mathbf{K} = \tilde{\mathbf{K}} + \tilde{\mathbf{K}}^\top$

nius norm analogs. Indeed, if a scenario were to arise in which the screening portions of the algorithm began to dominate the cost, the concept of shell blocks could be generalized further to arbitrary blocks and a multi-tiered prescreening approach could be used. Another difference in our implementation is that dense data structures were used to store all $O(N^2)$ quantities. The primary reason for this is that the SCF solver in `MPQC` is currently based on an $O(N^3)$ eigensolve; alternative $O(N)$ SCF solvers are well known (e.g., density matrix minimization[39]) but are not yet implemented in our program.

It is difficult to measure the practical scaling of an algorithm outside of the context of its implementation details and execution environment. Thus, many authors[40–43] opt to present their algorithmic scaling in terms of integral counts, contraction sizes, or other implementation- and execution-independent metrics. Here we use several such metrics corresponding to particular sections of Algorithm 4. Our implementation of CADF-LinK has been heavily optimized to run well on massively parallel computers and to take advantage of both thread and process parallelism. However, a thorough discussion of the challenges involved in the parallelization of CADF-LinK has been reserved for a separate paper[44] in the interest of both saving space and appealing to a more general audience.

For our cost metrics, we chose a series of one dimensional molecules (linear alkanes) and a series of three dimensional systems (water clusters). Cartesian coordinates are included in the supplemental information. We used the basis set pairs Def2-SVP[45]/Def2-SVP/JK,[46] cc-pVTZ[47]/cc-pVTZ/JK,[48] and cc-pVQZ/cc-pVQZ/JK. For these series of molecules and basis sets, we recorded the number of three-center integrals computed (line 5), the number of multiplies in the contraction to form $\mathbf{B}$ (line 11), and the number of multiplies in the contractions to form $\tilde{\mathbf{K}}$ (lines 16 and 21). These data were averaged over the first three SCF iterations; later iterations were omitted to minimize variability due to convergence acceleration procedures.

For the computations in the remainder of this work, an SCF convergence criterion of $10^{-6}$ was used. Our method has no difficulty converging further than this, but since the emphasis is on per iteration performance, we chose a relatively loose convergence threshold. An initial screening threshold $\epsilon_K$ of $10^{-6}$ was used throughout as well. However, this number was varied for differential density iterations as follows: the screening threshold for a given iteration was taken to be the initial threshold if the full density was used or the ratio of differential density to full density Frobenius norm times the initial threshold if a differential density was used, down to a minimum threshold of $10^{-11}$. For the SQV$\ell$ estimate, $10^{-1}$ was used for both $\vartheta_{\text{ws}}$ and $\vartheta_{\text{SQ}}$.[37] This use of a scaled screening threshold allowed us to use a quite aggressive initial threshold without sacrificing as much accuracy in the final result. This amounts to a simplified version of a previous variable precision SCF approach[49] that yields excellent results for large molecules. Indeed, we were usually able to converge SCF computations to a root mean squared density change of $10^{-10}$ or less using an initial threshold of $10^{-6}$. This thresholding scheme also spreads the workload across iterations in a fairly uniform manor: in almost all cases, no iteration costs more than a factor of 2 different from any other iteration (without threshold scaling, factors of 5-10 were often observed).

## V. RESULTS AND DISCUSSION

### A. Scaling with respect to cost metrics

To analyze the asymptotic scaling of a positive computational cost metric $C(N)$ with the system size parameter $N$ it is convenient to introduce of an effective scaling exponent:

$$k(N_i) \approx \log_{\frac{N_i}{N_{i-1}}} \frac{C(N_i)}{C(N_{i-1})}. \tag{30}$$

This amounts roughly to a "two-point fit" to the form $C(N) = aN^k + b$, for constants $a$ and $b$, at the points

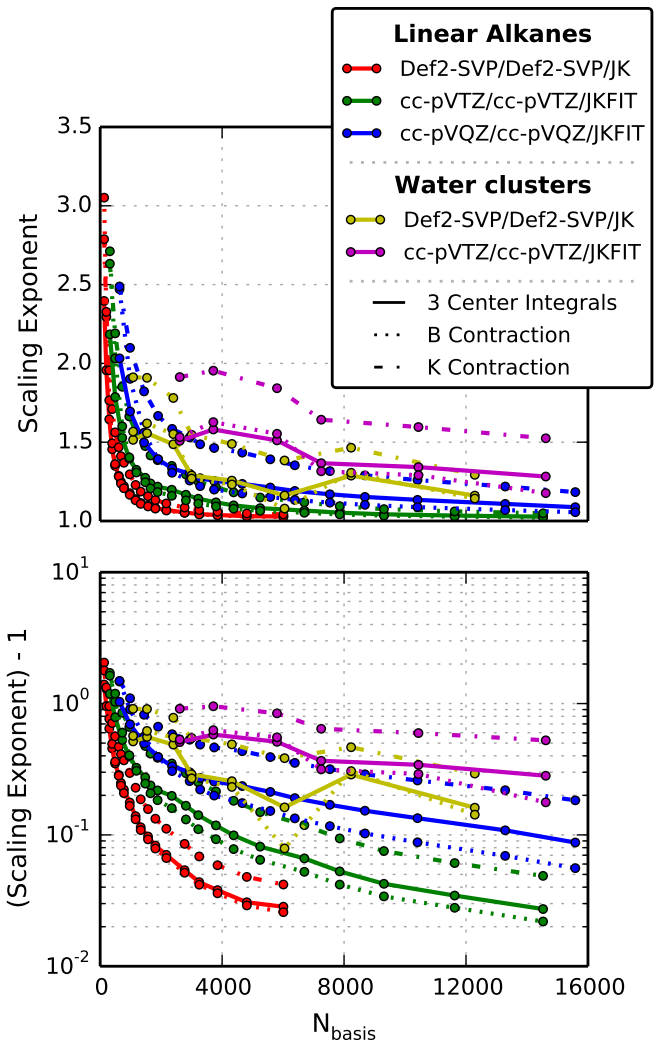


FIG. 1: Effective scaling exponents (Eq. (30)) of the three cost metrics of the CADF-LinK algorithm described in Section IV. The top plot uses a linear range, and the bottom plot shows the same data on a logarithmic range (subtracting 1 to show proximity to linear scaling).

$N_i$ and $N_{i-1}$. By definition, the effective exponent of the cost of a $O(N)$ algorithm should approach 1.0 in the limit of large $N$.

Figure 1 shows the effective scaling exponents for the key cost metrics of our algorithm when applied to our one- and three-dimensional series of molecules. It is clear that all scaling exponents are less than 2 for $N > 2000$, and decrease monotonically with $N$. Also, the **K** contraction step has the worst scaling (highest effective exponent). This is anticipated due to the relatively weak restriction of the loop in line 20 of Algorithm 4, which only utilizes Schwarz screening. However, since the non-LinK behavior of the **K** contraction is better than other parts of the algorithm ($O(N^3)$ with no screening compared to $O(N^4)$ for the **B** contraction without screening), the **K** contraction still costs less in terms of CPU time than the **B** contraction, even for our largest computations. Nevertheless, the scaling with respect to the **K** contraction metric is good, reaching $O(N^{1.1})$ around 2500 basis functions for linear alkanes with a small basis and $O(N^{1.2})$ around 4000 and 10000 basis functions for a triple zeta and quadruple zeta basis, respectively. Even for three dimensional systems, the scaling of the **K** contraction is still around $O(N^{1.5})$ by about 4000 and 15000 basis functions for the Def2-SVP/Def2-SVP/JK and cc-pVTZ/cc-pVTZ/JK basis sets, respectively. The scaling behavior of the more expensive sections, the integral computation and **B** contraction, is even better. The scaling reaches $O(N^{1.1})$ around 1300, 3500, and 8700 basis functions, respectively, for linear alkanes with the three basis set pairs in our study. For our three dimensional systems, the scaling of the **B** contraction is about $O(N^{1.5})$ by around 2400 and 5800 basis functions for Def2-SVP/Def2-SVP/JK and cc-pVTZ/cc-pVTZ/JK, respectively.

While the scaling exponent functions are less smooth for water clusters compared to linear alkanes due to the less systematic growth of the former, the overall trends in the data are similar for both one- and three-dimensional systems. Obviously, the decay of the scaling exponents is much more rapid for one dimensional systems and for smaller, less diffuse basis sets. In the worst case—three dimensional water clusters with the cc-pVTZ/cc-pVTZ/JK basis set—the scaling of the most computationally intense sections of the algorithm is around $O(N^{1.25})$ for a system with 756 atoms. Given that in the context of HF and hybrid KS DFT the typical basis sets are smaller and less diffuse than cc-pVTZ, it completely reasonable to conclude that CADF-LinK will closely approach linear scaling behavior in the vast majority of use cases.

### B. Errors

Figure 2 shows the errors in absolute energies resulting from the CADF-LinK approximation, relative to CADF-SCF with Schwarz screening only (the Coulomb matrix was computed in the same manner in both sets of computations). As with the CADF approximation and Hartree–Fock theory itself, the absolute energy errors arising from the CADF-LinK approximation scale linearly with the size of the system. However, these errors are still about one or two orders of magnitude smaller that the errors arising from the CADF approximation for molecules of comparable size (see Ref. 33), which are in turn smaller than the basis set errors or the errors of the Hartree–Fock method itself. Importantly, as demonstrated by the bottom plot in Figure 2, the absolute error per electron does not increase significantly beyond a point for the molecules we tested.

The larger error per electron for three dimensional structures compared to one dimensional ones can be heuristically rationalized as follows. As two basis func-

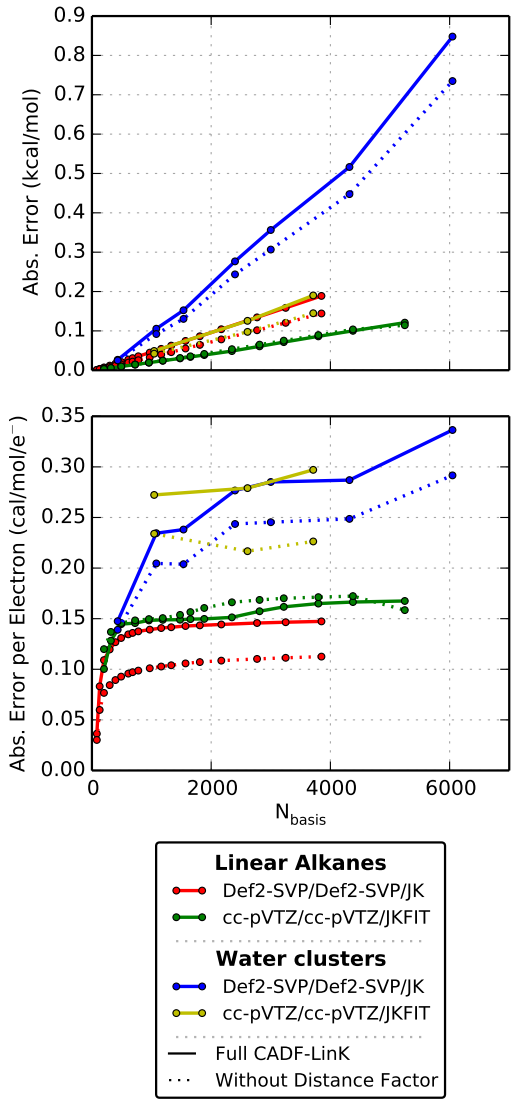


FIG. 2: Errors in absolute energies arising from the CADF-LinK approximation, relative to CADF-based exchange using only Schwarz screening. Errors per electron are given in the bottom plot to show that the CADF-LinK errors increase linearly with system size.

tions are separated from each other, their pair-wise contribution (roughly speaking) to the energy is excluded at some distance. A three dimensional structure will have more basis functions at or near this distance than a one dimensional structure. Furthermore, in the specific case of water clusters compared to linear alkanes, the latter have a significantly larger number of covalent bonds, meaning that a much smaller portion of the contributions to the energy are likely to have separations near this critical exclusion length.

Another notable feature of Figure 2 is the error behavior when the distance factor (from the SQV$\ell$ estimator) is omitted. As the original SQV$\ell$ paper noted,[37] the estimator performs worse for the Def2-$X$VP/Def2-$X$VP/JK basis sets than the cc-pV$X$Z/cc-pV$X$Z/JK basis sets, because the former contains more contracted basis functions, particularly in the fitting basis, than the latter. Since the SQV$\ell$ estimator is tightest for uncontracted integrals, the error arising from its use with contracted basis functions is expected to be larger than with basis sets containing fewer contracted functions. The data in Figure 2 show this, but they also show that the increase in error from the use of the SQV$\ell$ estimator is relatively small: in the Def2-SVP/Def2-SVP/JK case with linear alkanes, the SQV$\ell$ estimator causes roughly a factor of two increase in the CADF-LinK error, and for water clusters the increase is even less.

### C. Speedups

Figure 3 shows the speedups for the three cost metrics we examined relative to a CADF-based exchange build with Schwarz screening only—that is, the ratios of each cost metric for CADF exchange with Schwarz screening only to the cost metric for CADF-LinK for a given system size. The most substantial speedups—as much as 250 times—are seen in the **B** contraction, which is expected given that both the inner and outer loops are restricted by CADF-LinK (by the $L^{(B)}_{\mu\lambda}$ and $L^{(3)}_{\mu\lambda}$ lists, lines 10 and 4 respectively in Algorithm 4), whereas the 3 center integrals and the **K** contraction are only restricted by one list each. Indeed, the only significant additional restriction of the contraction in line 21 relative to Schwarz screening comes from the outermost loop over $(\mu_a X_c)$ pairs (since the stronger concentricity restriction of the loop in line 19 is imposed even without the formation of $L^{(B)}_{\mu X}$). Consequently, the least substantial speedups are seen in the **K** contraction, though again this section is the least computationally intense of the three even for the largest computations we performed. These data clearly show that the small increase in error (see Section V B) is, in most contexts, more than compensated for by a massive decrease in computational effort.

Another noteworthy trend in the data from Figure 3 is that for the first two cost metrics (3 center integrals and the **B** contraction), larger speedups were seen for three dimensional systems, while for the **K** contraction, larger speedups were seen for linear alkanes. This is attributed to the more prominent role of the distance factor in the $L^{(3)}_{\mu X}$ definition compared to the $L^{(B)}_{\mu X}$ definition. For $L^{(3)}_{\mu X}$, the distance factor $\tilde{R}^{X}_{\mu\lambda}$ contributes directly to the thresholding for the inclusion of $\lambda$, whereas for $L^{(B)}_{\mu X}$, the distance factors for all $\lambda$ connected by a P-junction to a given $\sigma$ contribute to the thresholding for the inclusion of $\sigma$. Heuristically, the 3 center integrals cost metric is most influenced by the efficiency of $L^{(3)}_{\mu X}$, and the **K** contribution is most influenced by the efficiency of $L^{(B)}_{\mu X}$ (the **B** cost metric is influenced by both). Since we have already noted in Section V B that the three-

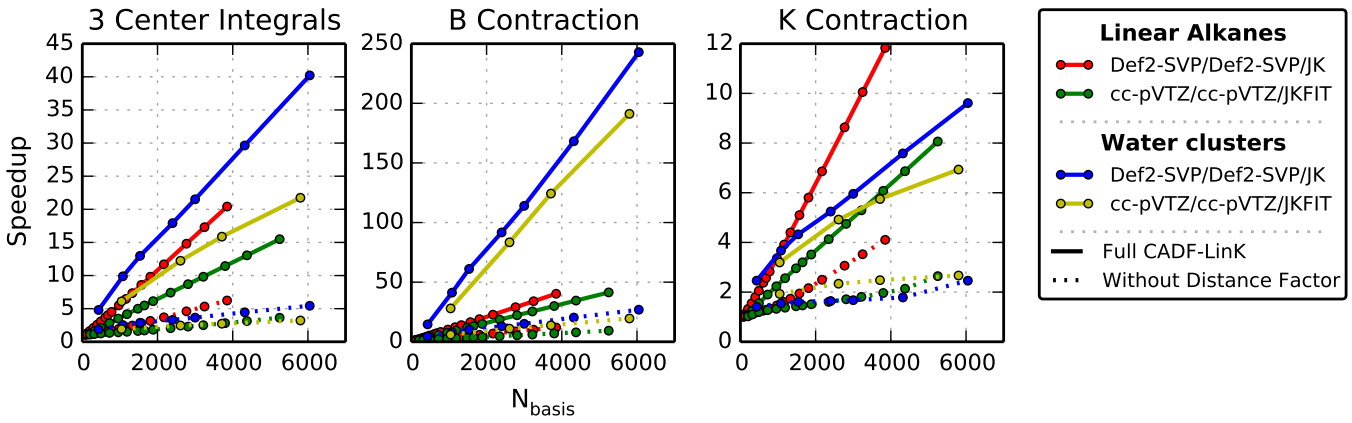


FIG. 3: CADF-LinK speedups for the three cost metrics examined here relative to a CADF-based exchange build with Schwarz screening only. The dotted lines show speedups for CADF-LinK without the distance-including SQV$\ell$ estimator.

dimensional structures are more susceptible to the effects of the distance-including estimator, it makes sense that more substantial gains would be seen with water clusters for the 3 center integrals and the **B** contraction and with linear alkanes for the **K** contraction.

The more interesting trend in Figure 3 is the contribution to the speedup from the distance-dependent integral screening. For large water clusters with the Def2-SVP/Def2-SVP/JK basis pair, nearly tenfold speedups are observed compared to CADF-LinK without the SQV$\ell$ estimator. This is remarkable, given that in the context of four-center ERI-based LinK the use of distance-dependent screening results in speedups of about 2.0.[41] This further demonstrates that there is more to be gained from screening three-center ERIs than from screening their four-center counterparts.

## VI. CONCLUSIONS

We have presented a linear scaling algorithm for computing the exchange matrix using the concentric atomic density fitting approximation. Our algorithm has been shown to perform well for basis sets of all sizes, and we have carried out some of the largest triple- and quadruple-zeta basis computations ever to demonstrate this point. Errors in absolute energies from the CADF-LinK approximation are substantially smaller than other sources of error and have been shown to grow linearly with basis set size. Even for large basis sets, our method shows near linear scaling for systems of less than 1000 atoms, and for smaller basis sets the onset of linear scaling is even more rapid. Not only does our algorithm serve as a highly efficient way to compute the exchange matrix in the context of Hartree-Fock and DFT methods, but also it offers a blueprint for the use of the concentric atomic density fitting in many-body methods.

## VII. ACKNOWLEDGEMENTS

The research by DSH and EFV was supported by NSF grants CHE-0847295, CHE-1362655, and ACI-1047696, and a Camille Dreyfus Teacher-Scholar Award. The research by DSH and HFS was supported by NSF grant CHE-1361178. This work used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U. S. Department of Energy under Contract No. DE-AC02-05CH11231.

[1] A. D. Becke, J. Chem. Phys. **98**, 1372 (1993).
[2] D. R. Hartree, Rep. Prog. Phys. **11**, 113 (1947).
[3] C. Riplinger and F. Neese, J. Chem. Phys. **138**, 034106 (2013).
[4] S. Goedecker, Phys. Rev. B **58**, 3501 (1998).
[5] J. Almlöf, "Modern electronic structure theory," (World Scientific, Singapore, 1995) p. 110.
[6] E. Schwegler and M. Challacombe, J. Chem. Phys. **105**, 2726 (1996).
[7] E. Schwegler, M. Challacombe, and M. Head-Gordon, J. Chem. Phys. **106**, 9708 (1997).
[8] J. C. Burant, G. E. Scuseria, and M. J. Frisch, J. Chem. Phys. **105**, 8969 (1996).
[9] C. Ochsenfeld, C. A. White, and M. Head-Gordon, J. Chem. Phys. **109**, 1663 (1998).
[10] E. Schwegler and M. Challacombe, J. Chem. Phys. **111**, 6223 (1999).
[11] R. A. Friesner, Chem. Phys. Lett. **116**, 39 (1985).
[12] F. Neese, F. Wennmohs, A. Hansen, and U. Becker, Chem. Phys. **356**, 98 (2009).
[13] E. G. Hohenstein, R. M. Parrish, and T. J. Martinez, J. Chem. Phys. **137**, 044103 (2012).
[14] R. M. Parrish, E. G. Hohenstein, T. J. Martinez, and C. D. Sherrill, J. Chem. Phys. **137**, 224106 (2012).
[15] E. G. Hohenstein, R. M. Parrish, C. D. Sherrill, and T. J. Martinez, J. Chem. Phys. **137**, 221101 (2012).
[16] A. Sodt and M. Head-Gordon, J. Chem. Phys. **128**, 104106 (2008).
[17] P.-O. Löwdin, J. Chem. Phys. **21**, 374 (1953).
[18] J. L. Whitten, J. Chem. Phys. **58**, 4496 (1973).
[19] E.-J. Baerends, D. E. Ellis, and P. Ros, Chem. Phys. **2**, 41 (1973).
[20] J. A. Jafri and J. L. Whitten, J. Chem. Phys. **61**, 2116 (1974).

[21]M. W. Feyereisen, G. Fitzgerald, and A. Komornicki, Chem. Phys. Lett. **208**, 359 (1993).
[22]O. Vahtras, J. Almlöf, and M. W. Feyereisen, Chem. Phys. Lett. **213**, 514 (1993).
[23]K. Eichkorn, O. Treutler, H. Öhm, M. Häser, and R. Ahlrichs, Chem. Phys. Lett. **240**, 283 (1995).
[24]K. Eichkorn, F. Weigend, O. Treutler, and R. Ahlrichs, Theor. Chim. Acta **97**, 119 (1997).
[25]R. Bauernschmitt, M. Häser, O. Treutler, and R. Ahlrichs, Chem. Phys. Lett. **264**, 573 (1997).
[26]F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs, Chem. Phys. Lett. **294**, 143 (1998).
[27]R. T. Gallant and A. St-Amant, Chem. Phys. Lett. **256**, 569 (1996).
[28]A. Sodt, J. E. Subotnik, and M. Head-Gordon, J. Chem. Phys. **125**, 194109 (2006).
[29]S. Reine, E. Tellgren, A. Krapp, T. Kjærgaard, T. Helgaker, B. Jansik, S. Host, and P. Salek, J. Chem. Phys. **129**, 104101 (2008).
[30]C. F. Guerra, J. G. Snijders, G. Te Velde, and E.-J. Baerends, Theor. Chem. Acc. **99**, 391 (1998).
[31]M. A. Watson, N. C. Handy, and A. J. Cohen, J. Chem. Phys. **119**, 6475 (2003).
[32]M. Krykunov, T. Ziegler, and E. v. Lenthe, Int. J. Quantum Chem. **109**, 1676 (2009).
[33]D. S. Hollman, H. F. Schaefer, and E. F. Valeev, J. Chem. Phys. **140**, 064109 (2014).
[34]P. Merlot, T. Kjærgaard, T. Helgaker, R. Lindh, F. Aquilante, S. Reine, and T. B. Pedersen, J. Comput. Chem. **34**, 1486 (2013).
[35]B. I. Dunlap, J. Mol. Struc. **529**, 37 (2000).
[36]M. Häser and R. Ahlrichs, J. Comput. Chem. **10**, 104 (1989).
[37]D. S. Hollman, H. F. Schaefer, and E. F. Valeev, J. Chem. Phys. (submitted).
[38]C. Janssen, E. Seidl, and M. Colvin, in *ACS Symposium Series, Parallel Computing in Computational Chemistry*, Vol. 592 (1995) pp. 47–61.
[39]M. Challacombe, J. Chem. Phys. **110**, 2332 (1999).
[40]M. Beer and C. Ochsenfeld, J. Chem. Phys. **128**, 221102 (2008).
[41]S. A. Maurer, D. S. Lambrecht, D. Flaig, and C. Ochsenfeld, J. Chem. Phys. **136**, 144107 (2012).
[42]S. A. Maurer, D. S. Lambrecht, J. Kussmann, and C. Ochsenfeld, J. Chem. Phys. **138**, 014101 (2013).
[43]G. E. Scuseria and P. Y. Ayala, J. Chem. Phys. **111**, 8330 (1999).
[44]D. S. Hollman, H. F. Schaefer, and E. F. Valeev, in preparation.
[45]F. Weigend and R. Ahlrichs, Phys. Chem. Chem. Phys. **7**, 3297 (2005).
[46]F. Weigend, Phys. Chem. Chem. Phys. **8**, 1057 (2006).
[47]T. H. Dunning, J. Chem. Phys. **90**, 1007 (1989).
[48]F. Weigend, Phys. Chem. Chem. Phys. **4**, 4285 (2002).
[49]N. Luehr, I. S. Ufimtsev, and T. J. Martinez, J. Chem. Theory Comput. **7**, 949 (2011).